

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
14 August 2003 (14.08.2003)

PCT

(10) International Publication Number
WO 03/067835 A1

(51) International Patent Classification⁷: H04L 12/56, H04J 3/16

(21) International Application Number: PCT/US03/02927

(22) International Filing Date: 31 January 2003 (31.01.2003)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/353,254 1 February 2002 (01.02.2002) US

(71) Applicant: TELLABS OPERATIONS, INC. [US/US];
1415 West Diehl Road, Naperville, IL 60563 (US).

(72) Inventors: CRANE, Benjamin, Mack, T.; 5616 Elinor
Avenue, Downers Grove, IL (US). TELLER, Susan,
W.; 1600 Mayapple Court, Naperville, IL 60565 (US).
SADLER, Jonathan, B.; 44 Waxwing Avenue, Naperville,
IL 60565 (US).

(74) Agent: FISH, Charles, S.; Baker & Botts, L.L.P., 2001
Ross Avenue, Suite 600, Dallas, TX 75201-2980 (US).

(81) Designated States (*national*): AE, AG, AL, AM, AT (utility model), AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ (utility model), CZ, DE (utility model), DE, DK (utility model), DK, DM, DZ, EC, EE (utility model), EE, ES, FI (utility model), FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SK (utility model), SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZM, ZW.


(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

- with international search report
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

[Continued on next page]

(54) Title: METHOD AND APPARATUS FOR MULTI-LAYER NETWORK IN SONET /SDH



(57) Abstract: Each node of a telecommunications network determines a connection type attributes available for each signal type supported by the node. Each signal type represents a different connection routing layer within the telecommunications network. Adaptation costs involved in traversing from one connection routing layer to another connection type attributes and adaptation costs are included in a link state advertisement broadcasted by each node in the telecommunications network. A route calculation is performed for a desired signal to determine a route through the telecommunications network for the signal. The route calculation takes into account adaptation costs in determining the shortest route for the signal through the telecommunications network.

BEST AVAILABLE COPY

WO 03/067835 A1



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

METHOD AND APPARATUS FOR MULTI-LAYER NETWORK IN SONET/SDH

TECHNICAL FIELD OF THE INVENTION

The present invention relates in general to telecommunications network control processing and more particularly to a method and system for multi-layer network routing.

BACKGROUND OF THE INVENTION

Calculation of a route through a network is based on link attributes advertised by each node in the network. There are several known link attributes that may be advertised by nodes within telecommunications network. These link attributes include traffic engineering metric, maximum or total reservable bandwidth, unreserved bandwidth, resource class/color, link protection type, and shared risk link group. The traffic engineering metric specifies the link metric for traffic engineering purposes. The maximum or total reservable bandwidth specifies the maximum bandwidth that may be reserved on this link in one direction. Unreserved bandwidth specifies the amount of bandwidth not yet reserved on the link in the one direction. Resource class/color specifies administrative group membership for this link. The link protection type specifies the protection capability that exists for the link. The shared risk link group attribute identifies a set of links that share a resource whose failure may affect all links in the set.

The link state advertisement may also include an Interface Switching Capability Descriptor. The Interface Switching Capability Descriptor describes the switching capability for an interface where the link is defined as being connected to a node by an interface. For example,

an interface that connects a given link to a node may not be able to switch individual packets but it may be able to switch channels within a Synchronous Optical Network (SONET) payload. Interfaces at each end of a link may not have the same switching capabilities. For bi-directional links, the switching capabilities of the link are defined to be the same in either direction for data entering and leaving the node through that interface. For a unidirectional link, it is assumed that the Interface Switching Capability Descriptor at the far end of the link is the same as at the near end of the link. A unidirectional link is required to have the same interface switching capabilities at both ends of the link.

The Interface Switching Capability Descriptor may specify a switching capability, an encoding type, a maximum and minimum labeled switch path (LSP) bandwidth, and interface maximum transmission unit. The switching capability descriptor specifies whether the interface is layer 2, packet, time division multiplex, lambda, or fiber capable as well as whether the interface supports more than one of these types. Maximum LSP bandwidth specifies the smaller of the unreserved bandwidth and the maximum reservable bandwidth per priority. Minimum LSP bandwidth specifies the minimum amount of bandwidth that may be reserved. The interface maximum transmission unit descriptor defines the maximum size of a packet that can be transmitted on this interface without being fragmented. Descriptors other than the switching capability descriptor are dependent on the type of switching capability defined in the switching capability descriptor.

5 The link attributes and interface switching capability descriptor discussed above are advertised by a node for its own egress interface only and require route calculations to find the reverse advertisement for a bi-directional link in order to determine the capabilities of a neighbor's end of the link. This unnecessarily complicates route calculation and adds the limitation that unidirectional links must have the same capabilities at both ends.

10 SUMMARY OF THE INVENTION

15 From the foregoing, it may be appreciated by those skilled in the art that a need has arisen for a technique to provide link state advertisements in a telecommunications network in order to facilitate multi-layer routing. In accordance with the present invention, a method and system for multi-layer network routing are provided that substantially eliminate or greatly reduce disadvantages and problems associated with conventional route calculation techniques.

20 According to an embodiment of the present invention, there is provided a method for multi-layer network routing that includes determining signal types implemented at each node of a network. Each signal type indicates a connection routing layer in the network. Connection types for each signal type and connection routing layer of each node of a network are determined for each link of each node. An availability of each connection type is also determined. The signal types, 25 connection types, and availability are broadcast from each node to each neighboring node in the network. Using this broadcasted information, a route through the 30

connection routing layers of the network may be determined for each signal sent in the network.

The present invention provides various technical advantages over conventional data management techniques. Some of these technical advantages are shown and described in the description of the present invention. Embodiments of the present invention may enjoy some, all, or none of these advantages. Other technical advantages may be readily apparent to one skilled in the art from the following figures, description, and claims.

BRIEF DESCRIPTION OF THE DRAWINGS

To provide a more complete understanding of the present invention and features and advantages thereof, reference is made to the following description, taken in conjunction with the accompanying figures, wherein like reference numerals represent like parts, in which:

FIGURE 1 illustrates a simplified diagram of a telecommunications network;

FIGURE 2 illustrates a simplified diagram of a multiple connection routing layer configuration of the telecommunications network;

FIGURE 3 illustrates connection type attributes determined and advertised at each node in the telecommunications network;

FIGURES 4A-C illustrate an example link state advertisement generated at each node;

FIGURES 5A-B illustrate a process flow for generating a route calculation implementing the connection type attributes;

FIGURE 6 illustrates an additional step involved in the generation of candidate nodes performed during the route calculation;

FIGURE 7 illustrates an additional step involved in the generation of candidate nodes during the route calculation;

FIGURE 8 illustrates a layer isolated connection routing approach to a multi-layer network.

DETAILED DESCRIPTION OF THE INVENTION

FIGURE 1 is a simplified diagram of a telecommunications network 10. Telecommunications network 10 includes a plurality of switching points or nodes 12 interconnected by links 14. Each node 12 is operable to transfer telecommunications signals using one or more signal types. Examples of signal types include Digital Service Level 1 (DS1); DS3, Virtual Tributary Level 1.5 (VT1.5), Synchronous Transport Signal Level 1 (STS-1), STS-3c, and Optical Carrier Level 3 (OC-3). Nodes 12 may also support other conventional signal types readily known by those skilled in the art. Each signal type represents a different connection routing layer within telecommunications network 10. For each signal carried within telecommunications network 10, a route to its intended destination is determined. Determination of the route through telecommunications network 10 may be made at an originating node 12, a node 12 acting as a supervisory or control node, or at a centralized management node 16 according to a desired design for telecommunications network 10. Determination of a route and generation of related information associated with nodes 12 of telecommunications network 10 may be performed by appropriate processing modules.

To automatically provision a signal through telecommunications network 10, a route is calculated and a connection setup is derived from the calculated route

for informing nodes 12. The route identifies the switching equipment at nodes 12 and the appropriate links 14 that the signal traverses in order to reach a destination node 12 from an originating node 12. Once the route has been calculated, signaling of the connection setup to the appropriate nodes 12 in telecommunications network 10 is performed in order to establish the intermediate switching connections necessary to provide end to end service for the signal. Alternatively, the connections can be configured from a central source or by other means according to desired implementations. Signaling for the connection setup is performed by standard techniques known by those skilled in the art.

FIGURE 2 shows an example layered view of telecommunications network 10. Routing of transport network signals over telecommunications network 10 may be accomplished by layer isolated connection routing or multi-layer connection routing. A layer is an abstraction containing the switching points and links that operate within one signal type. For the example shown in FIGURE 2, layer 20 is associated with the DS1 signal type, layer 22 is associated with the DS3 signal type, and layer 24 is associated with the STS-1 signal type. A client layer is considered to be the originally requested network layer and a server layer provides a trunk capability for the client layer. Other layers may also be present within telecommunications network 10 with one layer being present for each signal type supported by telecommunications network 10.

Layer isolated connection routing provides for a separate routing instance to exist for each layer supported by telecommunications network 10. These

routing instances operate independently of each other. Thus, each layer is unaware of the potential connectivity available within the other layers. By not knowing the connectivity available within other layers, the client layer is not able to optimally determine a point to request a server layer connection between two points in the client layer. In this manner, the client layer may pick non-optimal points for requesting a server layer connection in order to complete the route at the client layer. Moreover, the paths available at the server layer may not meet certain service requirements. If the client layer is unable to identify the characteristics of an available server layer connectivity in advance of making the request for a trunk connection, the client layer must wait for the route calculation to complete before determining whether the service requirements can be met. If the service requirements cannot be met, the client layer will need to reconsider the route it has taken and attempt to find another route that allows the service requirements to be met. Though the server layer may provide a list of connection possibilities with their respective attributes to the client layer, developing a list of paths and attributes at the client layer for each destination within the server layer would be processor and memory intensive when a network has a high number of paths between endpoints in the server layer and each path could have a different set of attributes. The present invention provides an embodiment to effectively use layer isolated connection routing in a multi-layer connection routing environment.

Multi-layer routing provides for a single routing instance to be responsible for connection routing within multiple layers of telecommunications network 10. This

routing instance is able to view the status and attributes of all links within multiple layers and allows for better routes to be determined with fewer processor cycles and less memory. Conventional link advertisement is insufficient to support route calculation involving adaptation to mapped and multiplexed server layers and does not allow for different costs to be assigned to link connections at different network layers on the same advertised link. There is typically not sufficient information to determine acceptable routes in multi-layer networks unless all boundary nodes are capable of switching at all network layers, which may not be possible in a network design. Also, current routing practices do not support routing through pooled resources for adaptation and/or interworking functions and do not address all types of routing constraints that are necessary for effective route calculation in multi-layer networks. The present invention provides an embodiment that turns these disadvantages into advantages for effectively providing multi-layer routing.

Routes are calculated through the use of an algorithm. A common algorithm used in the calculation of a route through a network is known as Dijkstra's algorithm. Dijkstra's algorithm is a standard subroutine for finding shortest paths from an originating node to a destination node taking into account different weights or costs involved in traversing through the network. A variation of Dijkstra's algorithm, referred to as extended Dijkstra's algorithm or Constrained Shortest Path First (CSPF) algorithm, takes into account a link's attributes, availability guarantees, and economic cost to use to determine if the link can satisfy the routing constraints provided in a routing request to a specified

destination node. In this manner, geographic diversity of a signal may be provided, economic cost for a connection may be minimized, or other desirable routing behaviors may be produced. In order to perform multi-layer network routing, each node determines link states and advertises these link states throughout telecommunications network 10. At each layer supported by telecommunications network 10, each node advertises various attributes to neighboring nodes. These link state advertisements are used in performing the route calculations using Dijkstra's algorithm or other route determination techniques.

FIGURE 3 shows connection type attributes that may be supported by a node and advertised to a neighboring node. These connection type attributes add efficiency and flexibility to route calculations not available with conventional attribute advertisements. Connection type attributes include Transit 30, Source 31, Sink 32, Dangling Egress 34, Dangling Ingress 35, Dangling Source 36, and Dangling Sink 37. Connection type Transit 30 represents an ability of the advertising and neighbor node to receive transport network signals from a node and forward the transport network signals to the next node in the route. Connection type Source 31 represents an ability of the advertising node to originate transport network signals at the current layer and an ability for a neighbor node to forward the transport network signals to the next node in the route. Connection type Sink 32 represents the ability of the advertising node to forward the transport network signals to the neighbor node for termination. Connection type Dangling Egress 34 represents the ability of the advertising node to adapt transport network signals from the current connection

routing layer to a server connection routing layer. Connection type Dangling Ingress 35 represents an ability of a neighbor node to receive transport network signals at the current connection routing layer from a server connection routing layer. Connection type Dangling Source 36 represents an ability of an advertising node to originate transport network signals at the current connection routing layer and adapt the originating transport network signals to a server connection routing layer. Connection type Dangling Sink 37 represents an ability of a neighbor node to receive transport network signals at the current connection routing layer from a server connection routing layer and terminate the current transport network signals.

At physical medium layers, dangling relationships are not possible since the medium is assumed to be connected to the neighbor node and physical media are not subject to adaptation to other server layers. At physical signal layers, it is possible that equipment inserted into the physical link between nodes has properties that are more restrictive than the switching equipment at the ends of the link. For example, two transparent cross-connects may be connected by a fiber span that includes a SONET repeater. This would limit the link to SONET clients while the switching equipment does not have this limitation. In this case, the limiting properties of the span may be projected into the switches in a way that accurately indicates the type of connectivity available. The limitation is adopted by the nodes as a characteristic of the advertised link though the switching equipment themselves are not the source of the limitation.

FIGURES 4A-C show an example of an advertisement performed by a node 12 according to the present invention. In FIGURE 4A, the attributes unreserved bandwidth and maximum LSP bandwidth are shown. These attributes may be advertised once per link. The conventional attribute maximum or total reservable bandwidth is not needed since oversubscription may be handled by simply adjusting the unreserved bandwidth attribute by an appropriate fraction of the bandwidth reservation for each link state protocol assigned to the link.

FIGURE 4B shows a connectivity attribute group (CAG) which may be advertised as needed. The CAG includes the signal type, the connection types supported and available, server signal type availability and associated adaptation cost, client signal type availability and associated adaptation cost, and server endpoint affinity. Each node determines the neighbor nodes with which it communicates during an identification phase of operation. A capability exchange phase is then performed so that a node can learn the various capabilities of each of its neighbor nodes. With this information, a node can build CAGs for each link with each neighbor node. Alternatively, these phases may be provisioned at node installation.

A CAG is formed and advertised for each signal type supported by the node. The signal type identifies the layer for this link being advertised. The connection type identifies one or more of the connection types shown in FIGURE 3. The server signal type availability indicates that the node has the capability to adapt the current signal type to an identified server layer and that source connectivity is available at that server

layer on this link. The adaptation cost identifies the cost involved in extending the dangling egress connectivity in order to traverse to the identified server layer. The client signal type availability indicates that the neighbor node has the capability to adapt the current signal type to an identified client layer and that dangling connectivity is available at the identified client layer on this link. The adaptation cost identifies the cost involved in extending the sink connectivity in order to traverse to the identified client layer. The server affinity endpoint field identifies a list of router identifiers that indicate which nodes should have preferential treatment for termination of a server trail that originates at this node while routing this signal type.

The connection types and associated availability are shown as bit fields with one bit position for each connectivity type defined above. These fields indicate whether each connectivity is supported and currently available for this link. Though shown in this manner, the link state advertisement may take on any form as desired for the communication of this information. If a connection type is advertised as currently available for a link, this indicates that at least one link connection of that type is available. To support setup of multiple co-routed connections, the advertised information could be extended to include the number of available connections of each type. However, the benefit of this extension may not be worth the resulting increase in database size as crankback may be used to address the relatively rare cases in which the advertised information is not sufficient to guarantee that the calculated route is acceptable.

FIGURE 4C shows additional attributes that may be advertised for each link. These attributes may be repeated as needed before the CAG to which they apply. These attributes include resource class/color, link protection type, and shared risk link group. These attributes relate to various routing constraints. If they are not included, the default values apply. Any of these attributes may be individually repeated with a new value in order to set that value for the CAGs that follow.

FIGURES 5A-B show an example process in determining a route calculation using the advertised CAG information. The process shows how Dijkstra's algorithm is changed to utilize the CAG information to accomplish route extension via multiplexed and mapped server trails. Though changes to Dijkstra's algorithm are shown, such changes are shown for example purposes only as the advertised CAG information may be implemented in other link state based route determination techniques and not limited to application to Dijkstra's algorithm. In the process shown, each node has an associated node identifier and a signal stack to uniquely identify the node. The signal stack is a stack of signal types that represent the current layers for the connection being routed at this node. As with all stacks, new values are added or pushed on the top of the stack and values are removed or popped from the top of the stack.

The calculation performed by the process of FIGURES 5A-B yields a set of intra-area routes associated with an area, Area A. A shortest path tree is determined by a node calculating the route using a specified node in the network topology as a root. The formation of the shortest path tree is performed in two stages. In the

first stage, only links between nodes in a single client layer are considered. In the second stage, links within one or more server layers are considered. At each iteration of the algorithm, there is a list of candidate nodes. Paths from the root to these candidate nodes have been determined but the shortest paths have yet to be defined. Paths to the candidate node that is closest to the root are guaranteed to be shortest. A path is said to be shorter if it has a smaller link state cost. The link state cost of a path is the sum of the costs of the path's constituent links as they exist in a layer network (i.e., CAG). Upon identification, a candidate node is added to the shortest path tree and removed from the candidate node list. Nodes adjacent to the candidate node added to the shortest path tree are examined for possible addition to the candidate node list and the shortest path tree. The algorithm continues to reiterate until the candidate node list is empty.

In FIGURES 5A-B, process flow begins at block 50 where the data structures of the algorithm are initialized and candidate nodes are cleared. A shortest path tree is initialized by adding first a node, for example node V, representing the root with a signal stack containing the signal requested to be routed. Subsequently, the shortest path tree is updated to include new nodes meeting the desired routing constraints. Area A's Transit Capability is set to False. At block 52, the link state advertisement of node V added to the shortest path tree is examined. Each link described by the link state advertisement provides signal types and costs to neighboring nodes. For each described link beginning at block 54, if the attributes of the link between node V and a neighboring node, for example node

W, do not satisfy the routing constraints requested, the examined link is skipped and the next link is examined. If routing constraints are satisfied, process flow proceeds to block 56 where node W's link state advertisement is examined. At block 57, if the link state advertisement for node W does not exist, has reached a maximum age, or does not include a link back to node V, then the next link in Node V's link state advertisement is analyzed. Otherwise, process flow proceeds to block 58 for an analysis of each CAG in the link. At block 59, if the attributes of the CAG do not satisfy the routing constraints, the next CAG in the link is analyzed. If routing constraints are satisfied here, process flow proceeds to block 60 where the signal type of the CAG is compared to the top of node V's signal stack. If there is no match, the next CAG in the link is examined.

Upon a match of a signal type, process flow proceeds to block 62 where the connection types of the CAG are examined in order to generate new candidate nodes. If transit connectivity is available, an instance of node W is formed with the current signal stack to specify uniqueness from one instance of node W to another. The advertised cost of each instance is set to the transit cost for the CAG. If sink connectivity is available, the signal stack for node V has more than one entry, and any client signal type of the CAG matches the second signal type in the signal stack of node V, an instance of node W, is formed for sink connectivity with the current signal stack excluding the top element of the signal stack for uniqueness and the advertised cost is set to the CAG adaptation cost for that client signal type. If dangling egress connectivity is available, an instance of node W

is formed with the current signal stack for dangling egress and for each server signal type in the CAG. The advertised cost is set to the CAG adaptation cost for each server signal type. For each instance of node W
5 generated, process flow proceeds to block 64 to determine if the newly generated instance of node W is already on the shortest path tree. If so, move on to the next generated instance of node W. If not, process flow proceeds to block 66 where the link state cost for the
10 path from the root to node W is calculated. The link state cost is the sum of the link state cost of the shortest path to node V and the advertised cost. At block 68, if the link state cost is greater than or equal to the value that already appears for node W in the
15 candidate node list, then the next CAG is examined. If the link state cost is less than the value that appears for node W on the candidate node list or if node W does not yet appear on the candidate node list, then an entry in the candidate node list for node W is set to the
20 calculated link state cost at block 69 and the next generated node W is examined.

Once all generated node W's for each CAG for each link have been examined, process flow proceeds to block 70 where a node in the candidate node list that is
25 closest to the root is chosen and added to the shortest path tree. At block 72, if the node added to the shortest path tree has a link state advertisement indicating that the destination is directly connected or reachable and the signal stack for the added node
30 contains one entry equal to the destination signal requested, then the calculation of the route is complete. The route is obtained by tracing back from this added node to the root of the shortest path tree. If the route

is not complete, process flow returns to block 52 to examine the link state advertisement for the newly added node.

When determining a route of a sub-network connection for a signal through a network, it is desirable to traverse from a client layer, where the signal originates, to a server layer and return. However, while the destination node may be accessible through both client and server layers, the destination node may not support the adaptation function needed to support the client signal on all of its interfaces at the server layer. As a result, the route calculation develops a path that uses an adaptation function to return the connection to the originating layer of the signal prior to completing the route. This behavior is supported through the sink connection type. Since the sink connection type has an associated signal type, it is possible to require the sink connection type match the type of signal being routed. In the case where a particular client signal type is not supported on an interface, the constraint matching function will invalidate the sink connection as a candidate for completing the signal routing. The route calculation will then continue to evaluate other candidates. This behavior is also preferable when the destination node for which a route is being calculated is not the terminating node but a border node which connects this route calculation domain to another domain. A link can be identified as being a border by looking at the advertisement type found in the supporting routing protocol.

When determining the route of a signal through telecommunications network 10, adaptation functions

advertised on a link become candidates for extending that route through a server layer. This candidate will be evaluated along with other candidates that are within the same layer as the signal being routed. The cost of the adaptation function, or adaptation cost, becomes a determining factor as to when a server layer connection will be necessary to complete the route for the signal. There may be times where it is desirable to control the points at which server layer extension will be allowed in route calculation. This does not change the fact that an adaptation function exists and it is not removed from the advertisement described above. Instead, a special adaptation cost, such as 0xffff, can be used to denote that route extension through a server layer is prohibited for the signal.

FIGURE 6 shows the change in the process of FIGURES 5A-B to implement this special adaptation cost. Since the adaptation cost is listed as part of the CAG associated with traversing a signal from a client layer to a server layer, the review of the adaptation cost occurs at the time that the adaptation is pushed onto the signal path. The availability of the adaptation also is checked when traversing from the server layer back to the client layer. Consequently, the return from server layer to client layer imposes a reverse check on the availability of the client layer to server layer adaptation function. If the reverse check fails during the route calculation, then the return to the client layer cannot be accomplished at this point and a new candidate will then be evaluated. The check for the special adaptation cost occurs during the generation of new candidate nodes. During the check for availability of a sink or dangling egress availability, if the

adaptation cost has the special adaptation cost value, in this case 0xffff, then a candidate node is not generated. If the adaptation cost is not the special adaptation cost value, then a candidate node is added to the candidate node list.

FIGURE 7 shows an example change to the process of FIGURES 5A-B implementing an affinity cost adjustment for route extensions via multiplexed server trails. When extending a route through a dangling egress followed by a server layer source connection, a special candidate route is created that has an added endpoint constraint that it must reach a matching server layer sink connection at a node that is already a neighbor of the current node for the link containing transit connections at the client layer being routed. This special candidate route has its cost adjusted by reducing or eliminating the adaptation cost for the server layer. The endpoint constraint is applied to this route candidate, and any route candidates that are generated from it, as the route calculation proceeds. The endpoint constraint for the special candidate route is determined by identifying all the neighbors of the current node on links that support transit connection types at the layer being routed regardless of whether or not these links currently have transit connections available. For the sink connectivity check of block 62, the added limitation on generating a node W is that the top of the signal stack for node V has either a NULL endpoint constraint or the endpoint constraint includes node W. For the dangling egress connectivity check, a node W is generated with both an endpoint constraint and a NULL endpoint constraint. The advertised cost for the first generated node W is set to the CAG transit cost and the advertised cost for the

second generated node W is set to the CAG adaptation cost. The endpoint constraint is a list of affinity endpoints that can be used to terminate the server layer trail and achieve the costs associated with the path to node V. Any signal in the signal stack, except the bottom signal, can be associated with a unique endpoint constraint. This method could be extended to create special candidate routes with different cost discounts depending on the number and size of existing trunks to each current neighbor node.

FIGURE 8 shows the implementation of layer isolated connection routing in a multi-layer network. Through the effective advertisement and route calculation described above, effective multi-layer network routing can be achieved. However, there may situations where it may be beneficial to provide layer isolated connection routing in a multi-layer network. A main limitation of layer isolated connection routing discussed above is the inability to take into consideration potential connectivity that can be supplied by server layer connections through links to the client layer in which the signal is to be routed. This limitation can be eliminated by assigning a node identifier to a fictitious node, or pseudo node 80, that serves to represent the potential connectivity provided by a server layer between switching equipment operating at the client layer.

Typically, a node identifier is assigned to a network node. The node identifier serves to correlate all the links advertised relative to that node as having a common switching point through which signals can be routed. In multi-layer networks, there may be resources that could be used to create additional connectivity or new links between nodes by creating new trails in a

server layer of the network. However, it would be inefficient to advertise all possible new links due to the vast number of possibilities. Instead, the potential connectivity provided by a server layer can be effectively summarized through the pseudo node. The pseudo node provides the possibility that any node at the edge of the server layer may be connected to any other node. The pseudo node represents the potential server layer connectivity in support of the client layer without the need to show the details of the server layer connectivity.

For each client layer node connected to a server layer, one or more links would be included in the link state database to represent this connectivity. These links would be included in the link state advertisement of the node in the client layer. Route calculations could then take into account the possibility of creating new paths in the server layer by traversing links to and from the pseudo node. The cost of these links can be set according to the desired policy regarding the preference that should be given to using existing links in the client layer versus the newly created links to the pseudo node using the server layer.

Advertisement of links from the pseudo node must also be made. Generally each node in the network advertises its outgoing links. However, the pseudo node is fictitious. To overcome this, the content of a link advertisement may be independent from its source. In distributed flooding protocols, the neighbor node from whom a node receives a link advertisement is often not the node originating that link advertisement. By exploiting this feature of flooding protocols, a client layer node is allowed to advertise a link from itself to

the pseudo node and also advertise the corresponding link from the pseudo node to itself. This advertisement is created as if it was generated by the pseudo node and then flooded throughout the network as any other link advertisement. This allows distributed network implementations to employ pseudo nodes without the added complexity of creating a separate pseudo node routing protocol entity.

Though pseudo nodes have been used in the past, the present invention is able to apply pseudo nodes technique to multi-layer transport networks. Pseudo node information being advertised does not require designated router election and other coordination between proxy advertisers. The use of pseudo nodes in a multi-layer network provides the ability to control, via provisioning, the communities connected to a given pseudo node to effect routing policies. Also, an ability is provided for nodes advertising connection to a pseudo node to recognize the pseudo node in routes and request a connection through the core network to replace pseudo node hops in the route.

In summary, effective multi-layer network routing can be performed through advertisement of appropriate link state information and through the use of the link state information during route calculation. The link state information includes connection type attributes that specify not only how a node can transport information in a layer and between layers of a network but also identify optimal points to move between layers and what nodes provide access to a desired layer of the network. Multi-layer network routing may also be represented in a layer isolated connection routing approach through the use of pseudo nodes.

The techniques performed by the present invention may be implemented in software, hardware, or a combination of both. For example, each node may have individual modules that can identify the different signal types and connection routing layers associated with the node, determine the connection types and availabilities corresponding to each connection routing layer at the node, and broadcast the link state advertisement, either in separate modules or different functions may be combined into the same module. Modules may also be provided to calculate a route through the network and determine the various transit and adaptation costs associated with connections in the network. These modules may provide the structure within the network for performing the functionality of the present invention.

Although the present invention has been described in detail with reference to particular embodiments, it should be understood that various other changes, substitutions, and alterations may be made hereto without departing from the spirit and scope of the present invention. For example, although the present invention has been described with reference to Dijkstra's algorithm, other routing calculations may be used with equal effectiveness in the present invention. Moreover, a number of potentially suitable components that facilitate the processing of information in various types of formats, any suitable objects, elements, hardware, or software may be used in the applications or operations described above. The arrangements described above in conjunction with telecommunications system 10 provide only an example configuration used for purposes of teaching and substitutions and modifications may be made where appropriate and according to particular needs.

Other changes, substitutions, variations, alterations, and modifications may be ascertained by those skilled in the art and it is intended that the present invention encompass all such changes, substitutions, variations, alterations, and modifications as falling within the spirit and scope of the appended claims. Moreover, the present invention is not intended to be limited in any way by any statement in the specification that is not otherwise reflected in the appended claims.

WHAT IS CLAIMED IS:

1. A method for multi-layer network routing, comprising:

5 determining signal types implemented at each node of a network, each signal type indicating a connection routing layer in the network;

determining connection types for each signal type and connection routing layer at each node of a network at each link of each node;

10 determining availability of each connection type;

broadcasting signal types, connection types, and availability from each node to each neighboring node in the network.

15 2. The method of Claim 1, further comprising:

including limiting properties of each link associated with each node in the connection type determination.

20 3. The method of Claim 1, further comprising:

calculating a route from an originating node to a destination node through different connection routing layers of the network.

25 4. The method of Claim 1, further comprising:

establishing a first adaptation cost at each node capable of providing a connection from a first connection routing layer to a second connection routing layer.

5 5. The method of Claim 4, further comprising:
 assigning a particular value to the first adaptation
 cost in order to prevent the connection from the first
 connection routing layer to the second connection routing
 layer.

10 6. The method of Claim 4, further comprising:
 establishing a second adaptation cost at each node
 capable of providing a connection from the second
 connection routing layer to the first connection routing
 layer.

15 7. The method of Claim 6, further comprising:
 assigning a particular value to the second
 adaptation cost in order to prevent the connection from
 the second connection routing layer to the first
 connection routing layer.

20 8. The method of Claim 1, further comprising:
 identifying nodes having preferential treatment for
 termination of a connection routing layer.

25 9. The method of Claim 1, further comprising:
 assigning a transit cost associated with each
 routing layer supported by a connection from one node to
 another.

30 10. The method of Claim 1, further comprising:
 providing information from a particular node as to
 its capability to traverse from a first connection
 routing layer to a second connection routing layer and a
 capability of a neighbor node to traverse from the second

connection routing layer to the first connection routing layer.

11. A network for communicating transport signals, comprising:

5 a plurality of nodes, the plurality of nodes operable to communicate transport signals over a plurality of layers of the network, each layer representing a different transport signal type where the originating layer is a client layer and the other layers are server layers, each node operable to generate and broadcast a link state advertisement, the link state
10 advertisement operable to indicate a connection capability of a particular node and a connection capability of a neighbor node to the particular node for each layer, the link state advertisement operable to determine through which layers in the network a transport
15 signal can be routed.

12. The network of Claim 11, wherein the connection capability of the particular node and the neighbor node are provided in a connection type field of the link state
20 advertisement, the connection type field operable to indicate any of a transit, source, sink, dangling egress, dangling ingress, dangling source, and dangling sink connection types associated with a link of the particular node.

25 13. The network of Claim 11, wherein the link state advertisement includes an availability and adaptation cost associated with traversing from one layer corresponding to another layer within the particular
30 node.

14. The network of Claim 11, wherein the link state advertisement includes an availability and an adaptation

cost associated with traversing from a server layer of the network to a client layer at the neighbor node.

- 5 15. The method of Claim 11, wherein the link state advertisement includes a list of nodes in the network having priority for terminating a trail in a server layer.

16. A computer readable medium including code for multi-layer network routing, the code operable to:

determine signal types implemented at each node of a network, each signal type indicating a connection routing layer in the network;

determine connection types for each signal type and connection routing layer at each node of a network at each link of each node;

determine availability of each connection type;

broadcast signal types, connection types, and availability from each node to each neighboring node in the network.

17. The computer readable medium of Claim 16, wherein the code is further operable to:

calculate a route from an originating node to a destination node through different connection routing layers of the network.

18. The computer readable medium of Claim 16, wherein the code is further operable to:

determine cost values associated with traversing from one connection routing layer to another.

19. The computer readable medium of Claim 16, wherein the code is further operable to:

identify a connection capability of a neighbor node;

provide the connection capability of the neighbor node in the broadcast from a particular node.

20. The computer readable medium of Claim 16, wherein the code is further operable to:

provide an indication to prevent traversing from one connection routing layer to another at a particular node.

21. A system for multi-layer network routing, comprising:

means for determining signal types implemented at each node of a network, each signal type indicating a connection routing layer in a network;

means for determining connection types for each signal type and connection routing layer at each node of a network at each link of each node;

means for determining availability of each connection type;

means for broadcasting signal types, connection types, and availability in a link state advertisement from each node to each neighboring node in the network.

22. The system of Claim 21, further comprising:

means for calculating a route through the network in response to the link state advertisement.

23. The system of Claim 22, further comprising:

means for providing transit cost associated with communicating a transport signal from one node to another;


means for providing an adaptation cost associated with communicating the transport signal from one connection routing layer to another connection routing layer.

24. The system of Claim 23, further comprising:

means for setting a special adaptation cost at a particular node, the special adaptation cost indicating that communicating the transport signal from one connection routing layer to another connection routing layer is not to be performed at the particular node.

25. The system of Claim 21, wherein the link state advertisement broadcast by a particular node of the network includes a connection type field, the connection type field operable to indicate any of a transit, source, sink, dangling egress, dangling ingress, dangling source, and dangling sink connection types associated with a link of the particular node.

FIG. 1






FIG. 2

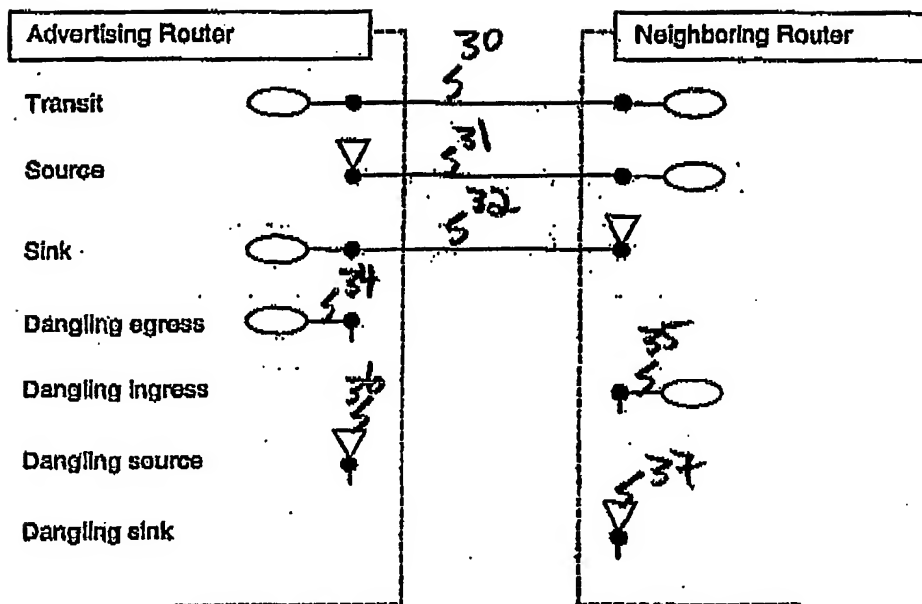


FIG. 3

[illegible]

76. 4A

1										2										3											
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1
Signal Type										Conn. Types										Availability											
Transit Cost																															
Number of Server Types															Number of Client Types																
Server Signal Type															Adaptation Cost																
Client Signal Type															Adaptation Cost																
Server Endpoint Affinity																															

FIG. 4B

U										L										A										J									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
Resource Class/Color																																							

```

0      1      2      3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|Link Prot. Type|                                     Reserved|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

```

0                                     1                                     2                                     3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Shared Risk Link Group                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Shared Risk Link Group                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

FIG. 4C

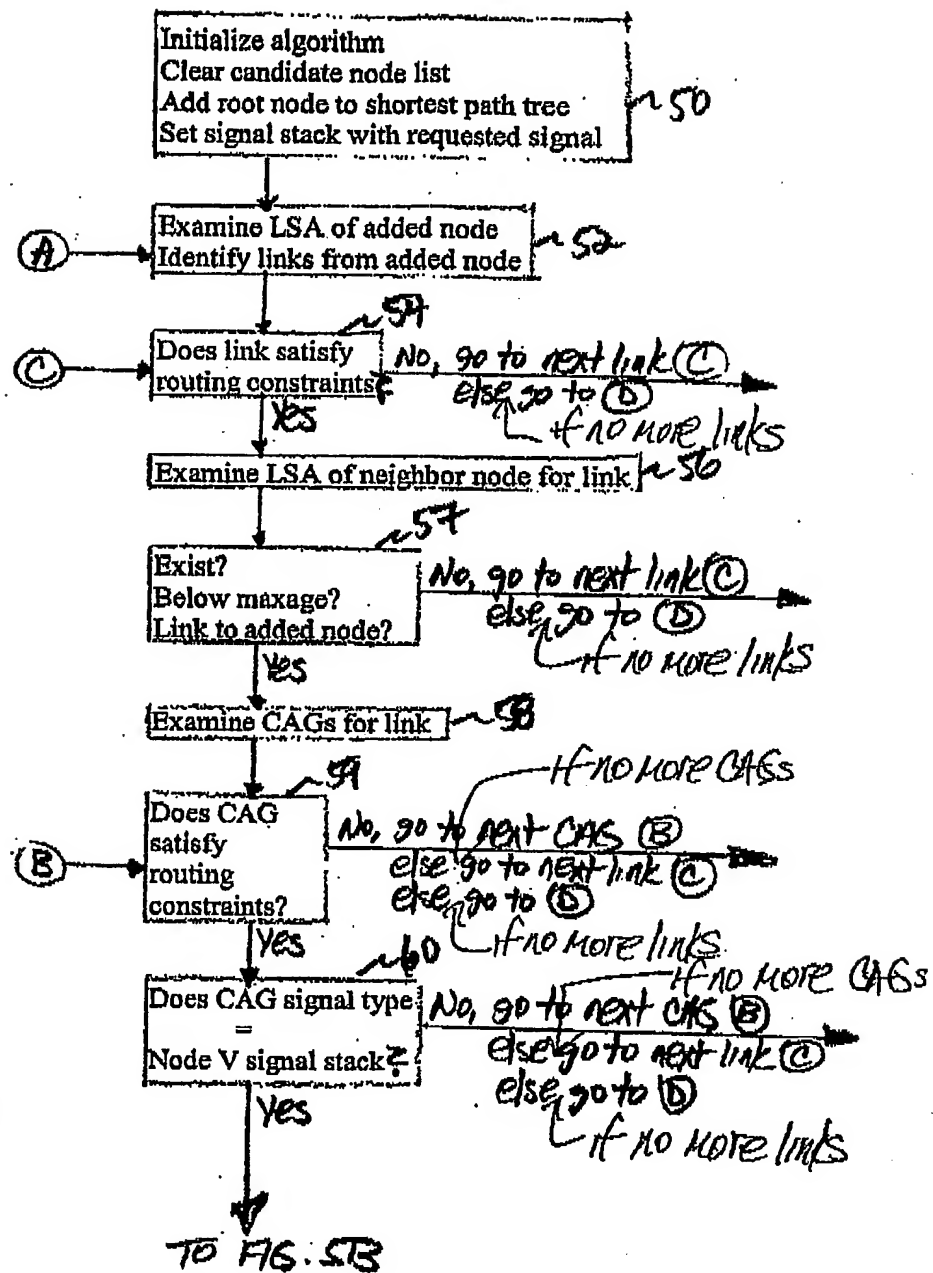
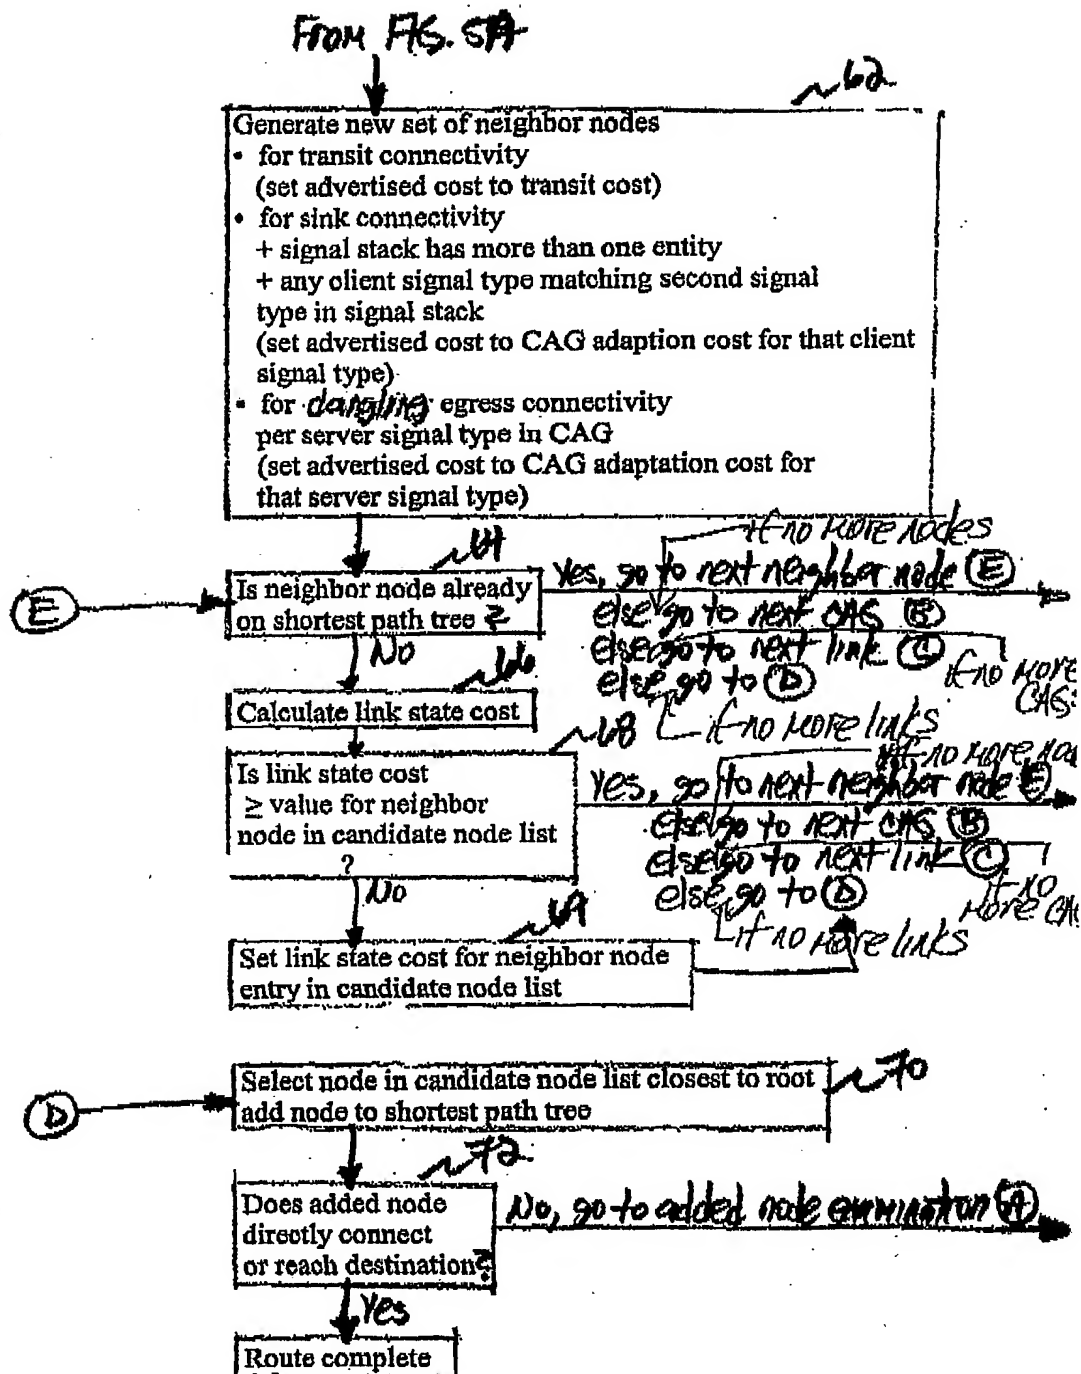


FIG. 5A



76. 5B

FROM FIG. 5A



n/62

- Generate new set of neighbor nodes
- for transit connectivity
(set advertised cost to transit cost)
 - for sink connectivity
 - + signal stack has more than one entity
 - + any client signal type matching second signal type in signal stack
 - + CAG adaptation cost not equal to 0xffff
(set advertised cost to CAG adaptation cost for that client signal type)
 - for dangling egress connectivity
per server signal type in CAG
(set advertised cost to CAG adaptation cost for that server signal type)



FIG. 6

FROM FIG. 5A

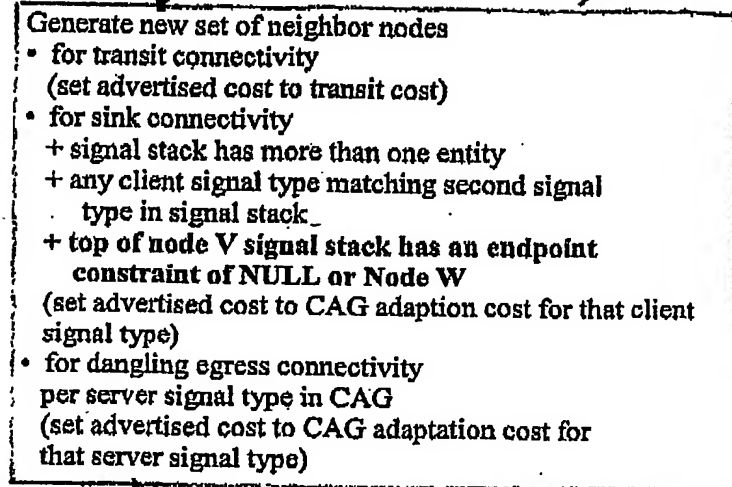


FIG. 7




FIG. 8

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 03/02927

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L12/56 H04J3/16

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 H04L H04J

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>US 2001/033548 A1 (ZADIKIAN HAIG MICHAEL ET AL) 25 October 2001 (2001-10-25)</p> <p>abstract paragraph '0026! - paragraph '0028! paragraph '0054! - paragraph '0056! paragraph '0076! paragraph '0126! paragraph '0135! - paragraph '0143! paragraph '0225!</p> <p>--- -/--</p>	<p>1-3, 8-11, 16-18, 21-23</p>

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *G* document member of the same patent family

Date of the actual completion of the international search

4 June 2003

Date of mailing of the international search report

12/06/2003

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Ngao, Y.S.

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 03/02927

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>EP 1 146 682 A (NORTEL NETWORKS LTD) 17 October 2001 (2001-10-17)</p> <p>paragraph '0042! - paragraph '0044! paragraph '0060! paragraph '0065! - paragraph '0066! paragraph '0072! - paragraph '0083! paragraph '0095! - paragraph '0100!</p>	<p>1-3, 8-11, 16-18, 21-23</p>
A	<p>EP 1 152 631 A (NORTEL NETWORKS LTD) 7 November 2001 (2001-11-07)</p> <p>paragraph '0022! - paragraph '0023! paragraph '0037! - paragraph '0040!</p>	<p>1-3, 11, 16-18, 21-23</p>

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 03/02927

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 2001033548	A1	25-10-2001	US 2003058804 A1	27-03-2003
			AU 3582000 A	01-08-2000
			WO 0042746 A1	20-07-2000
			US 2003031127 A1	13-02-2003
			US 2001048660 A1	06-12-2001
			US 2002054572 A1	09-05-2002
EP 1146682	A	17-10-2001	CA 2339716 A1	29-09-2001
			EP 1146682 A2	17-10-2001
EP 1152631	A	07-11-2001	CA 2343576 A1	04-11-2001
			EP 1152631 A2	07-11-2001
			JP 2002026990 A	25-01-2002

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant:

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☐ FADED TEXT OR DRAWING
- ☐ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☒ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☐ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.